

Present and Future Computing Requirements for:
**Simulations of and Model Validation on Small
Scale Plasma Experiments**

Vyacheslav (Slava) Lukin
Naval Research Laboratory

NERSC FES Requirements for 2017
March 19-20, 2013
Rockville, MD

NERSC Accounts Considered in this Case Study

Plasma Science and Innovation Center (PSI-Center):

PI: Thomas Jarboe, NERSC PI: Brian Nelson (U. of Washington)

- U. of Washington – Jarboe *et al.*
- U. of Wisconsin – Sovinec *et al.*
- Utah State U. – Held *et al.*
- Naval Research Lab – Lukin

Steady Inductive Helicity Injected Torus (HIT-SI):

PI: Thomas Jarboe, NERSC PI: Brian Nelson (U. of Washington)

Swarthmore Spheromak Experiment (SSX):

PI: Michael Brown (Swarthmore C.), NERSC PI: Vyacheslav Lukin (NRL)

- Swarthmore C. – Brown *et al.*
- Naval Research Lab – Lukin

Scientific Objectives



- Small scale magnetized plasma experiments usually focus on investigating one or a few specific phenomena that are known to be of concern, and have to be better understood and modeled, for the much more complex ITER-scale plasma experiments to operate reliably
- They often serve as test-beds for new ideas towards improving the design of a future fusion reactor
- With increased availability of computing resources and software, as well as the greater recognition of the intrinsic value of numerical modeling for understanding the behavior of highly nonlinear systems, advanced simulations have become an integral part of the research process on many of the small scale plasma experiments
- MHD based, particle based, and/or hybrid numerical plasma models developed by the greater computational plasma physics community are often employed for both idealized and whole-device simulations of these experiments.

Scientific Objectives

- The PSI-Center has been serving as the conduit to enable and support modeling efforts on many of the small and medium scale plasma experiments in the US, such as HIT-SI, SSX, Pegasus, ZaP etc., and, as of recently, internationally on MAST
- In the past, the PSI-Center's mission has been that of a support center to help the small experiments understand and interpret their data, as well as to help design new experiments
- Looking towards 2017, the mission of the PSI-Center will broaden to proactively perform systematic plasma model validation studies using the experimental data
- The goal is to take advantage of the small experiments' *in situ* and integral diagnostics with high rep rate at low cost to quantitatively estimate the validity of the plasma models implemented in the PSI-Center codes in the parameter regimes accessible to the experiments

Example: HIT-SI Simulations (Ackay, et al.)

HIT-SI experiment forms and sustains spheromaks with Steady Inductive Helicity Injection (SIHI) provided by two semi-toroidal injectors, X and Y.



NIMROD has been used to model the experiment

Example: HIT-SI Simulations (Ackay, et al.)

Hall MHD (hMHD) model with uniform density and $\beta = 0$


Experimental and Simulation Parameters ($T_i = T_e = 10$ eV)

	Simulation	Experiment
Injector flux (mWb)	1.2	0.6-1.4
Injector current (kA)	20	10-20
λ_{inj} (m ⁻¹)	20	15-25
f_{inj} (kHz)	14.5	14.5
$\langle n_e \rangle$ (m ⁻³)	1.5×10^{19}	$1-3 \times 10^{19}$
S	$3 \times 10^3 - 2 \times 10^4$	$10^3 - 10^4$
Pm	10-50	10-100?

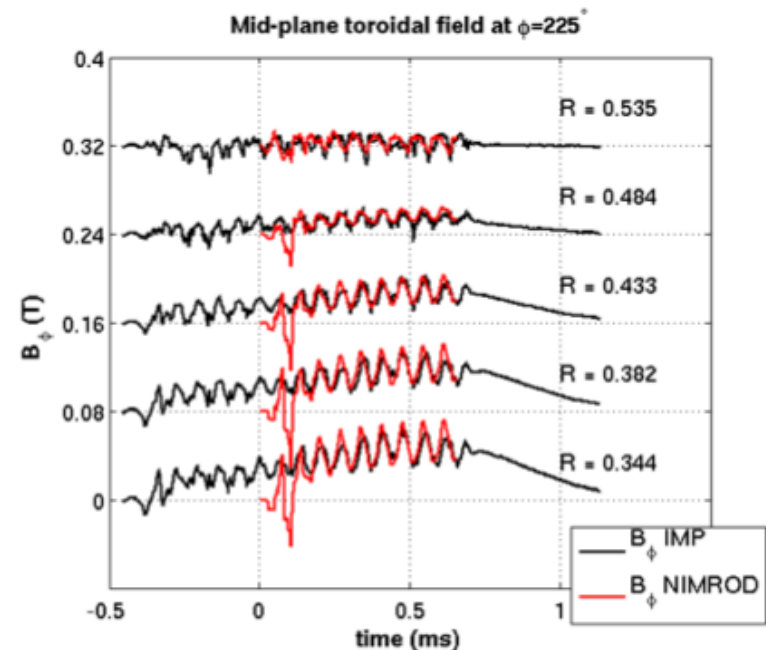
- $d_i = 8$ cm. $\Delta x = 0.5$ cm and $\Delta t = 3 \times 10^{-8}$ s.

Example: HIT-SI Simulations (Ackay, et al.)

Comparison of measured and calculated midplane internal magnetic fields




(a) B_z vs time (ms)



(b) B_ϕ vs time (ms)


Example: SSX Simulations (Lukin, *et al.*)



The spheromaks are injected from the coaxial magnetized plasma guns from either end of the machine into the flux conserver. The flux conserver is illustrated by the dark line. Representative flux surfaces of the two initial spheromaks are also shown.

Example: SSX Simulations (Lukin, *et al.*)

SSX diagnostics and parameters



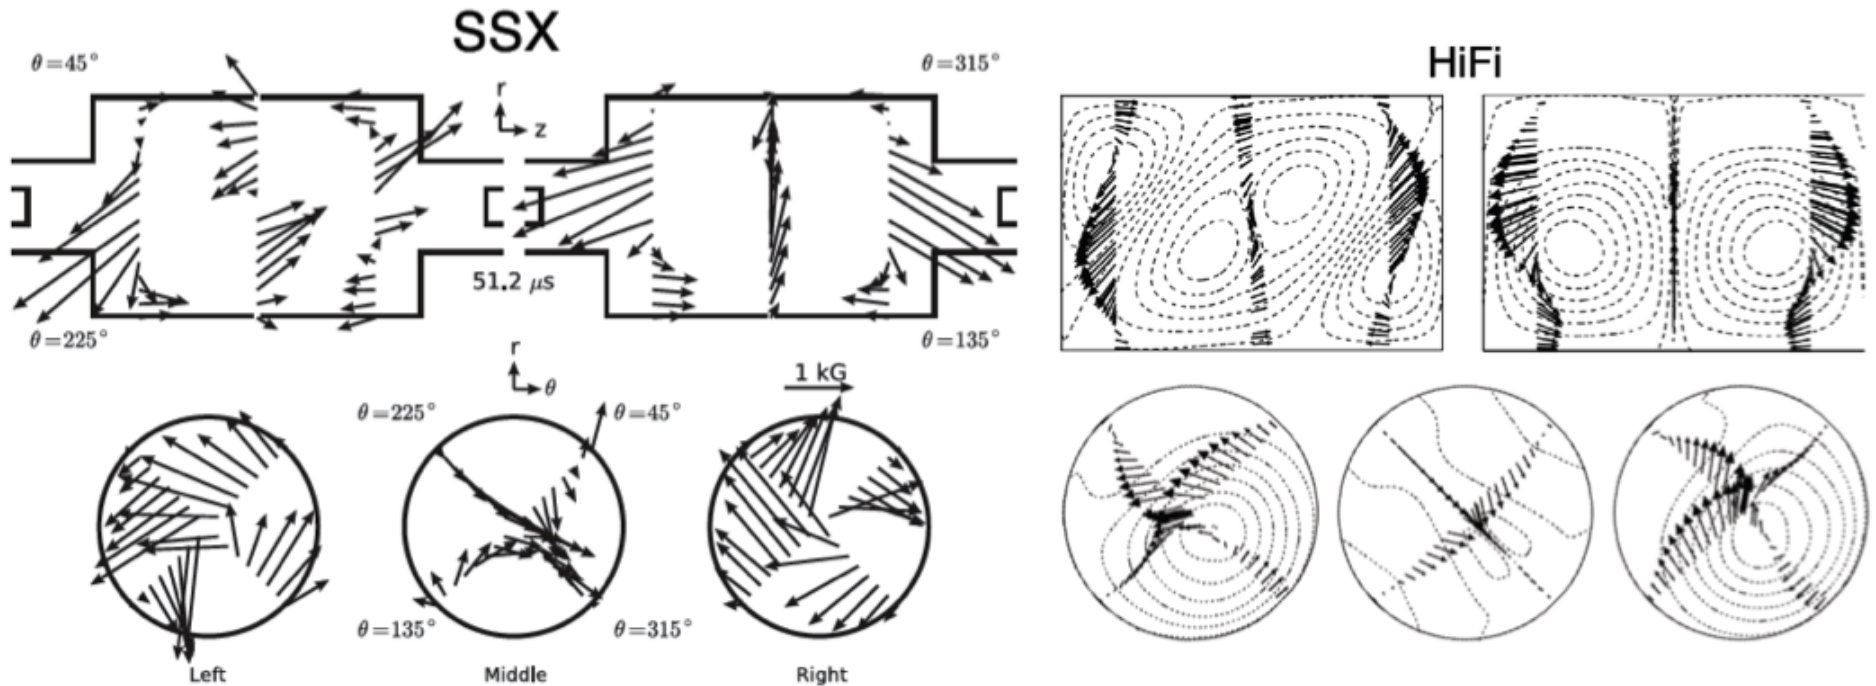
SSX plasma parameters

n_e	$10^{14} - 10^{15} \text{ cm}^{-3}$
T_e, T_i	20, 80 eV
B	1 kG
ρ_i	0.5cm
V_A	100 km/s
s	~ 10

- Magnetic probes inserted into the plasma
 - Equilibrium reconstructions by EQLFE CT equilibrium code
- $n_{e,avg}$ by HeNe interferometer
- T_i and flow measured by Ion Doppler Spectroscopy (IDS) system
- T_e measured by VUV monochrometer (CIV/CIII)


Example: SSX Simulations (Lukin, et al.)

3D resistive and Hall MHD simulations have been performed



Comparison of B-field measurements in equivalent planes from the SSX experiment and the HiFi simulation during the fast relaxation through null-point magnetic reconnection. The contours in the HiFi data represent in-plane B.

Example: MAST Simulations (Stanier, *et al.*)



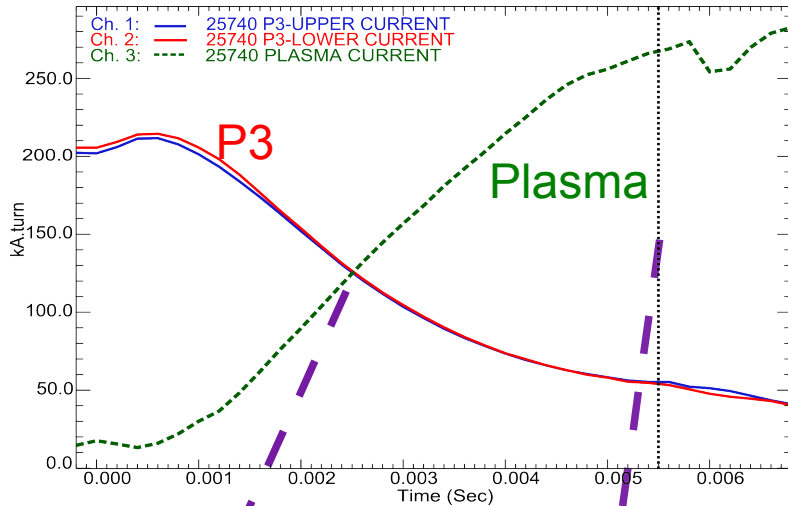


Normal Operating Parameters

- ▶ Major radius: $R = 0.85 \text{ m}$
 - ▶ Minor radius: $a = 0.65 \text{ m}$
 - ▶ $R/a = 1.3$

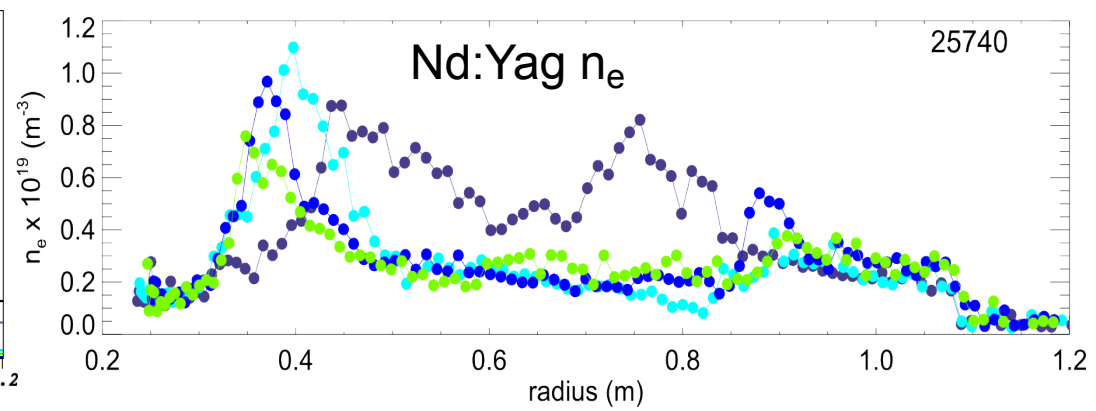
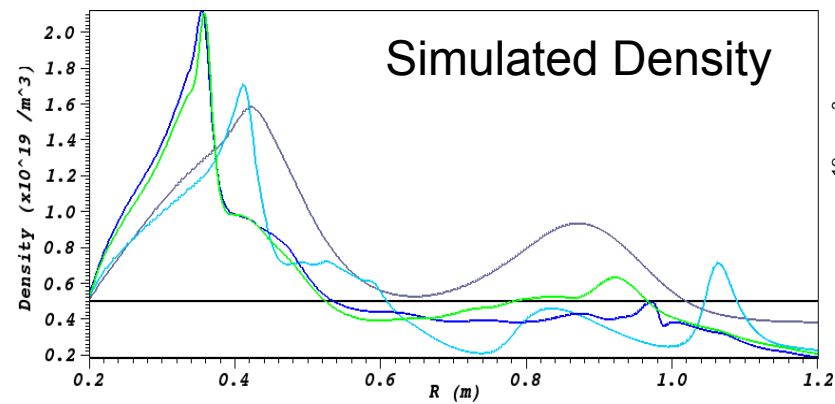
 - ▶ Toroidal field (at R): $B_T = 0.5 \text{ T}$
 - ▶ Current: $\leq 1.6 \text{ MA}$
 - ▶ Temperature $\sim 0.1 - 3 \text{ keV}$
 - ▶ Density: $10^{18} - 10^{20} \text{ m}^{-3}$
 - ▶ Ion Species: Deuterium
-
- ▶ P3 coils: used for **merging-compression start-up**
 - ▶ P4, P5: vertical field, P6: vertical position

Example: MAST Simulations (Stanier, et al.)

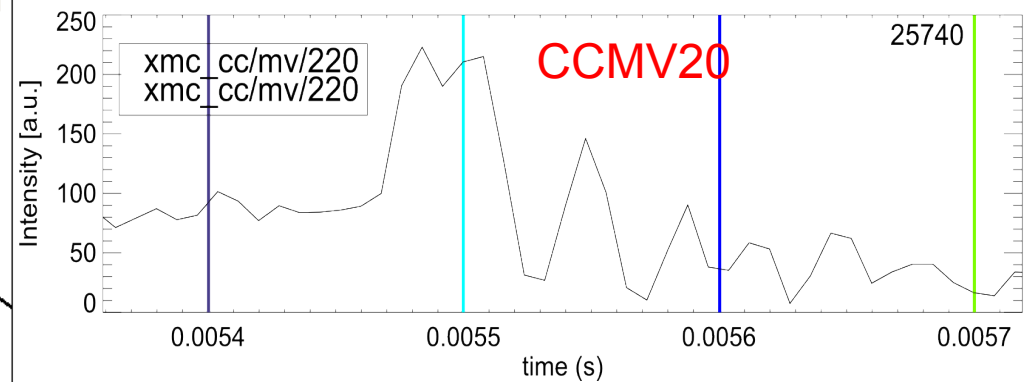
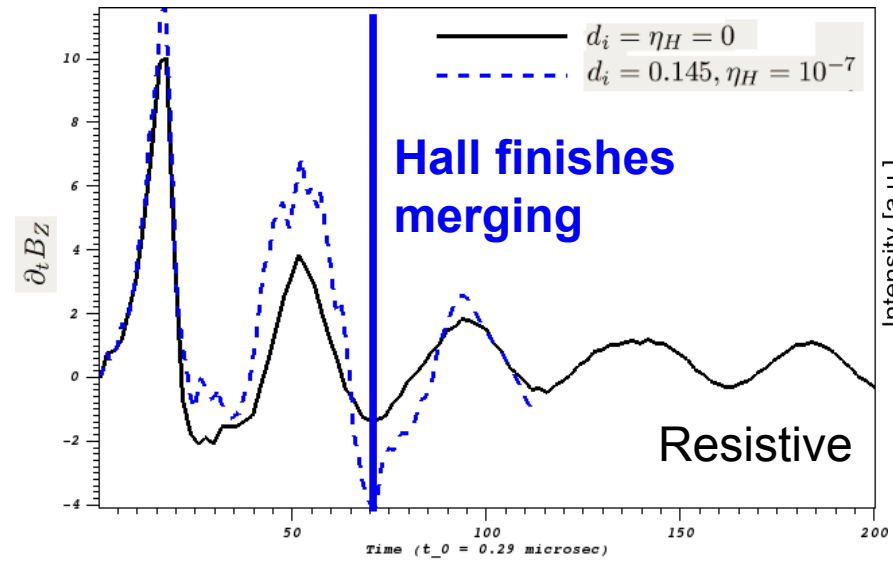


- ▶ Direct Induction (DI) start-up is expensive.
- ▶ Merging-Compression is an attractive alternative and is routinely used in MAST.
- ▶ Breakdown and induction around **P3 coils**.
- ▶ Merging via **reconnection** at mid-plane to form single ST.
- ▶ Compression via ramp-up of vertical field.
- ▶ Final state $I_{plasma} \propto I_{P3}$
- ▶ Up to **0.5 MA plasma current** obtained.
- ▶ Up to $T_e = 1$ keV achieved in on **ms timescale**.

Example: MAST Simulations (Stanier, et al.)



► Double peaked density profile. Evolution similar to Nd:Yag profiles.



► Simulated Mirnov is too fast (by factor of 3): other PF coils may be important.

Computational Strategies

- To conduct validation studies, need to perform many medium-sized simulations that complement the “one-off” biggest possible simulations by exploring the available parameter space and performing statistical studies around the performance point of any given experiment
- By far the most computationally intensive calculations presently performed and supported by the PSI-Center are two- and three-dimensional initial value extended or multi-fluid MHD calculations performed using the HiFi framework and the NIMROD code
- A typical calculation is either driven to or is initialized in an unstable state which in turn drives the nonlinear plasma dynamics of the experiment being simulated
- Other numerical codes: 3D force-free plasma equilibrium solver on an unstructured tetrahedral mesh PSI-Tet; DCON MHD stability code for axisymmetric toroidal plasmas; the particle trajectory-tracing code PUSH.

Computational Strategies

- Both NIMROD and HiFi are implicit initial value codes that employ semi-structured meshes and can use thousands of computing cores
- Both codes use semi-structured arbitrary order finite (spectral) elements spatial discretization in a two-dimensional plane; NIMROD uses the Fourier representation in the third dimension, HiFi treats all three dimensions on equal footing
- For grid generation, HiFi relies on the semi-structured grid generator package CUBIT, NIMROD uses a home-grown grid-generation algorithm
- Temporal advance: HiFi is fully implicit, i.e. the full system of PDEs is always advanced in a single implicit time-step – together with the generic user interface gives HiFi the flexibility of a general PDE solver; NIMROD uses a split-implicit time advance algorithm with hard-wired PDEs – allows for a smaller memory foot print

Computational Strategies

- Both NIMROD and HiFi use preconditioned Krylov-space methods for solving the resulting large sparse algebraic systems: NIMROD solver is customized with an interface to SuperLU_DIST; HiFi relies primarily on the PETSc interface for access to a variety of preconditioning and linear solver libraries
- Therein lies the biggest computational challenge and the limitation to parallel scaling...
- Strong scaling is limited to $O(100)$ cores
- Weak scaling is problem-dependent; reasonable scaling up to $O(10^4)$ cores
- No paradigm-changing shifts in the computational approach are expected by 2017; weak scalability improvement by a factor of 2 to 4 can be expected from algorithmic and linear solver library improvements

Present NERSC (and other) Usage

- We rely on the NERSC computational resources to a large extent: mostly use Hopper, also Copper, and Euclid for serial post-processing and visualization. Usage on the order of $(2-5) \times 10^6$ CPU hours
- However, local small (up to ~ 200 core) computing clusters also play a very important role in day-to-day research. Many users run local 25 to 100 core jobs, both in the code testing and the production simulation phases, on a daily basis
- Size of production runs can vary by about two orders of magnitude, from ~ 50 up to about 5000 cores, depending on the plasma parameter regime of the particular experiment being modeled, whether the simulation is done in two or three dimensions, and how detailed is the physical model used in the simulation

Present NERSC (and other) Usage

- There are two primary reasons for often using fewer cores than one could:
 - CPU hours is a limited resource and we want to use them efficiently. Due to lack of strong scaling beyond a couple of hundred cores, if certain amount of spatial resolution is sufficient and the problem fits on fewer cores and can run in reasonable time, using fewer cores is advantageous.
 - Much of the work we do is focused on understanding the underlying nonlinear physical processes that are at play in the experiments we model. The typical process involves many iterations with sensitivity studies to both numerical and physical model-dependent parameter variations. When for these purposes a smaller problem size is acceptable, due to the lack of strong scaling and limited available resources, using fewer cores is advantageous.

Present NERSC Memory Usage

- **Scratch (temporary) space:** ~700 GB

Used for storing running simulation checkpoint files, synthetic data analysis, and visualization.

- **Permanent:** ~2-5 GB

Used for storing libraries, source code and executables.

- **HPSS permanent archival storage:** ~2-5 TB

Used for data backup, and archiving past calculations for long term storage

- **I/O:**

NIMROD: serial with a single processor responsible for reading (writing) binary checkpoint files and appropriately distributing (collecting) the information

HiFi: both serial and parallel hdf5, as well as the serial binary format, are available as run-time I/O options for reading and writing checkpoint files

Projected 2017 HPC Requirements

Projected 2017 HPC Requirements

- **Uncertainty in the future domestic DoE FES budget as a whole, and the funding for the small scale plasma experiments, in particular, makes it difficult to estimate the computational hours that may be needed by PSI-Center and the collaborating experiments in 2017.**
- **With that said,** assuming a sufficient number of research scientists and graduate students is available to do the work, the expected need will be in the range of 100-200 Million CPU hours.

Projected 2017 HPC Requirements

- The primary reason for the increased need will be driven by the combination of two factors:
 - 1) Change in the mission towards model validation. The broader scope will require running many more simulations to quantify the uncertainty associated with the spatial resolution, initial and boundary conditions. Similarly, the sensitivity of the results to the specific physics model and the adjustable parameters used in the simulations will need to be evaluated. This alone will account for up to **100 fold** increase in the demand for computing time.
 - 2) The increased maturity and better performance of the codes will allow to devote more time to running simulations of the experiments and validating the models. Improvements in the codes' user interface will enable more users directly affiliated with the experiments to run the codes themselves. These may account for another **factor of 2-3** increase in the demand for NERSC compute hours.

Projected 2017 HPC Requirements

Memory, Data and I/O

- 2 GB/core memory if OK, 4 GB/core – much better!
- The permanent storage capacity is expected to scale linearly with the number of users, thus increasing by about a **factor of two**.
- The scratch space capacity is expected to scale linearly with the number of simulations being performed simultaneously and the size of the computational grids, thus likely increasing by a **factor of O(10)**.
- The HPSS archival storage is expected to increase linearly with the number of simulations performed over the course of a year, thus likely increasing by a **factor of O(100)**.
- No I/O bandwidth limitations for either HiFi or NIMROD have been identified so far.

Projected 2017 HPC Requirements

Support and Multi-Core Architectures

- Software Tools: PETSc, HDF5, NetCDF, SuperLU_DIST, MUMPS, HYPER, VisIt, Fortran-90 compilers
- Support: Maintaining high standards in account support, timely installation of new software versions, as well as continued training and tutoring in the use of new systems is all that is necessary
- Multi-Core: We now rely on MPI communication and run on homogeneous CPU systems. Initial development and planning for using heterogeneous many-core architectures has begun, but any help from the computer science community and NERSC will be welcome. (GPU-capable version of the PETSc library is now available, however GPU-enabled versions of such linear solvers as SuperLU_DIST and MUMPS, or similar alternatives, would have to become available for either code to be able to take full advantage of the new heterogeneous architectures.)

Final Comments

- Small scale plasma experiments are critical to the FES mission and their operation increasingly involves advanced modeling
- Validation of computational models on small scale experiments will become an important focus in the future
- This, together with an increased user base and **conditional on availability of funding for humans to run the codes**, is expected to lead to drastic increase in the number, but not the size of the simulations
- Thus, being limited by the weak scaling, increased throughput for 100 to 5000 core size jobs, possibly also allowing for longer runtimes for a single job, would be very helpful